



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/826,959	04/15/2004	James D. Ballew	064747.1015	1484
45507	7590	10/20/2008	EXAMINER	
BAKER BOTTS LLP 2001 ROSS AVENUE 6TH FLOOR DALLAS, TX 75201-2980				MEHRMANESH, ELMIRA
ART UNIT		PAPER NUMBER		
2113				
			NOTIFICATION DATE	
			DELIVERY MODE	
			10/20/2008	
			ELECTRONIC	

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

PTOmail3@bakerbotts.com
PTOmail4@bakerbotts.com

Office Action Summary	Application No.	Applicant(s)	
	10/826,959	BALLEW ET AL.	
	Examiner	Art Unit	
	Elmira Mehrmanesh	2113	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 14 July 2008.
 2a) This action is **FINAL**. 2b) This action is non-final.
 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1,4-11,14-21 and 24-31 is/are pending in the application.
 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
 5) Claim(s) 31 is/are allowed.
 6) Claim(s) 1,4-11,14-21 and 24-30 is/are rejected.
 7) Claim(s) _____ is/are objected to.
 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.
 10) The drawing(s) filed on 15 April 2004 is/are: a) accepted or b) objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
 3) Information Disclosure Statement(s) (PTO/SB/08)
 Paper No(s)/Mail Date 1/30/08, 3/14/08, 3/26/08, 10/1/08.

4) Interview Summary (PTO-413)
 Paper No(s)/Mail Date. _____.
 5) Notice of Informal Patent Application
 6) Other: _____.

DETAILED ACTION

This action is in response to an amendment filed on July 14, 2008 for the application of Ballew et al., for a "System and method for detecting and managing HPC node failure" filed April 15, 2004.

Information disclosed and listed on PTO 1449 has been considered.

Claims 1, 4-11, 14-21, 24-31 are pending in the application.

Claims 1, 11, 21, and 31 have been amended.

Claims 2-3, 12-13, and 22-23 have been cancelled.

Claims 1-30 are rejected under 35 USC § 103.

Claim 31 is allowed.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

The factual inquiries set forth in *Graham v. John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

Claims 1-30 are rejected under 35 U.S.C. 103(a) as being unpatentable over Huang (U.S. Patent No. 5,748,882) in view of Karpoff (U.S. PGPub No. 20010049740).

As per claim 1, Huang discloses a method comprising:

determining that one of a plurality of nodes has failed (col. 5, lines 15-19)

removing the failed node from a virtual list of nodes, the virtual list comprising one logical entry for each of the plurality of nodes (col. 10, lines 45-50)

determining that at least a portion of a job was being executed on the failed node (col. 7, lines 49-55) and terminating at least the portion of the job (Fig. 5, element 511).

determining that the job was associated with a subset of the plurality of nodes; and deallocating the subset of nodes from the job (col. 7, lines 5-67 through col. 8, lines 1-9)

Huang discloses an integrated fabric (col. 4, lines 66-67 through col. 5, lines 1-2). Each node contains communication links, communication ports (col. 10, lines 65-67), (col. 11, lines 60-61), and the fault tolerance socket (col. 18, lines 28-31). However Huang fails to explicitly disclose a switching fabric integrated onto a board and one or more processors.

Karpoff teaches:

each node comprising a switching fabric integrated to a card and at least two processors integrated to the card (FIG. 4A, a typical INFINIBAND® Architecture 20 includes one or more Central Processing Units (CPUs) 30, a Memory Controller 28, a Host Interconnect 29, a Host Channel Adapter (HCA) 22, a Target Channel Adapter (TCA) 24, and one or more Switches 26 (page 6, paragraph [0089]). Karpoff further

discloses nodes attached to the fabric can be assembled into logical subsets or partitions in order to group hosts or devices with like attributes, much like zoning capabilities of Fiber Channel fabrics (page 6, paragraph [0087]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to use the method and system for providing multimedia information of Karpoff in combination with the method for fault-tolerant computing of Huang to effectively monitor a multi-node system.

One of ordinary skill in the art at the time of the invention would have been motivated to make the combination because Huang discloses an integrated fabric (col. 4, lines 66-67 through col. 5, lines 1-2) wherein each node contains communication links, communication ports (col. 10, lines 65-67), (col. 11, lines 60-61). Huang's figure 2 shows an example of a node, which has at least one processor (col. 4, lines 66-67 through col. 5, lines 1-2). Karpoff's figure 4A shows a switching fabric, which includes one or more Central Processing Units (page 6, paragraph [0089]).

As per claim 4, Huang discloses each entry of the virtual list comprising a node status and the method further comprising changing the status of each of the subset of nodes to "available" (col. 10, lines 50-56).

As per claim 5, Huang discloses determining dimensions of the terminated job based on one or more job parameters and an associated policy; dynamically allocating

a second subset of the plurality of nodes to the terminated job based on the determined dimensions (col. 17, lines 1-21)
executing the terminated job on the allocated second subset (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 6, Huang discloses the second subset comprising a substantially similar set of nodes to the first subset (Fig. 2).

As per claim 7, Huang discloses dynamically allocating the second subset comprises: determining an optimum subset of nodes from a topology of unallocated nodes; and allocating the optimum subset (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 8, Huang discloses locating a replacement node for the failed node; and updating the logical entry of the failed node with information on the replacement node (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 9, Huang discloses determining that one of the plurality of nodes has failed comprises determining that a repeating communication has not been received from the failed node (col. 17, lines 21-30).

As per claim 10, Huang discloses determining that one of the plurality of nodes has failed comprises determining through polling that one of the plurality of nodes has

failed (col. 8, lines 43-63).

As per claim 11, Huang discloses software encoded in one or more computer-readable tangible media and when executed operable to:

determine that one of the plurality of nodes has failed (col. 5, lines 15-19)

remove the failed node from a virtual list of nodes, the virtual list comprising one logical entry for each of the plurality of nodes (col. 10, lines 45-50)

determine that at least a portion of a job was being executed on the failed node (col. 7, lines 49-55) and terminating at least a portion of the job (Fig. 5, element 511)

determine that the job was associated with a subset of the plurality of nodes; and deallocate the subset of nodes from the job (col. 7, lines 5-67 through col. 8, lines 1-9)

Huang discloses an integrated fabric (col. 4, lines 66-67 through col. 5, lines 1-2).

Each node contains communication links, communication ports (col. 10, lines 65-67), (col. 11, lines 60-61), and the fault tolerance socket (col. 18, lines 28-31). However Huang fails to explicitly disclose a switching fabric integrated onto a board and one or more processors.

Karpoff teaches:

each node comprising a switching fabric integrated to a card and at least two processors integrated to the card (FIG. 4A, a typical INFINIBAND® Architecture 20 includes one or more Central Processing Units (CPUs) 30, a Memory Controller 28, a Host Interconnect 29, a Host Channel Adapter (HCA) 22, a Target Channel Adapter (TCA) 24, and one or more Switches 26 (page 6, paragraph [0089]). Karpoff further

discloses nodes attached to the fabric can be assembled into logical subsets or partitions in order to group hosts or devices with like attributes, much like zoning capabilities of Fiber Channel fabrics (page 6, paragraph [0087]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to use the method and system for providing multimedia information of Karpoff in combination with the method for fault-tolerant computing of Huang to effectively monitor a multi-node system.

One of ordinary skill in the art at the time of the invention would have been motivated to make the combination because Huang discloses an integrated fabric (col. 4, lines 66-67 through col. 5, lines 1-2) wherein each node contains communication links, communication ports (col. 10, lines 65-67), (col. 11, lines 60-61). Huang's figure 2 shows an example of a node, which has at least one processor (col. 4, lines 66-67 through col. 5, lines 1-2). Karpoff's figure 4A shows a switching fabric, which includes one or more Central Processing Units (page 6, paragraph [0089]).

As per claim 14, Huang discloses each entry of the virtual list comprising a node status and the software further operable to change the status of each of the subset of nodes to "available" (col. 10, lines 50-56).

As per claim 15, Huang discloses to determine dimensions of the terminated job based on one or more job parameters and an associated policy; dynamically allocate a

second subset of the plurality of nodes to the terminated job based on the determined dimensions (col. 17, lines 1-21)

executing the terminated job on the allocated second subset (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 16, Huang discloses the second subset comprising a substantially similar set of nodes to the first subset (Fig. 2).

As per claim 17, Huang discloses the software is operable to dynamically allocate the second subset comprises software operable to: determine an optimum subset of nodes from a topology of unallocated nodes; and allocate the optimum subset (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 18, Huang discloses to locate a replacement node for the failed node; and update the logical entry of the failed node with information on the replacement node (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 19, Huang discloses the software being operable to determine that one of the plurality of nodes has failed comprises software operable to determine that a repeating communication has not been received from the failed node (col. 17, lines 21-30).

As per claim 20, Huang discloses the software being operable to determine that one of the plurality of nodes has failed comprises the software being operable to determine through polling that one of the plurality of nodes has failed (col. 8, lines 43-63).

As per claim 21, Huang discloses a system comprising:

- a plurality of nodes (Fig. 2)
- a management node (Fig. 2, element 104) operable to:
 - determine that one of the plurality of nodes has failed (col. 5, lines 15-19)
 - remove the failed node from a virtual list of nodes, the virtual list comprising one logical entry for each of the plurality of nodes (col. 10, lines 45-50)
 - determine that at least a portion of a job was being executed on the failed node (col. 7, lines 49-55) and terminating at least a portion of the job (Fig. 5, element 511)
 - determine that the job was associated with a subset of the plurality of nodes; and deallocate the subset of nodes from the job (col. 7, lines 5-67 through col. 8, lines 1-9)
- Huang discloses an integrated fabric (col. 4, lines 66-67 through col. 5, lines 1-2). Each node contains communication links, communication ports (col. 10, lines 65-67), (col. 11, lines 60-61), and the fault tolerance socket (col. 18, lines 28-31). However Huang fails to explicitly disclose a switching fabric integrated onto a board and one or more processors.

Karpoff teaches:

each node comprising a switching fabric integrated to a card and at least two processors integrated to the card (FIG. 4A, a typical INFINIBAND® Architecture 20 includes one or more Central Processing Units (CPUs) 30, a Memory Controller 28, a Host Interconnect 29, a Host Channel Adapter (HCA) 22, a Target Channel Adapter (TCA) 24, and one or more Switches 26 (page 6, paragraph [0089]). Karpoff further discloses nodes attached to the fabric can be assembled into logical subsets or partitions in order to group hosts or devices with like attributes, much like zoning capabilities of Fiber Channel fabrics (page 6, paragraph [0087]).

It would have been obvious to one of ordinary skill in the art at the time of the invention to use the method and system for providing multimedia information of Karpoff in combination with the method for fault-tolerant computing of Huang to effectively monitor a multi-node system.

One of ordinary skill in the art at the time of the invention would have been motivated to make the combination because Huang discloses an integrated fabric (col. 4, lines 66-67 through col. 5, lines 1-2) wherein each node contains communication links, communication ports (col. 10, lines 65-67), (col. 11, lines 60-61). Huang's figure 2 shows an example of a node, which has at least one processor (col. 4, lines 66-67 through col. 5, lines 1-2). Karpoff's figure 4A shows a switching fabric, which includes one or more Central Processing Units (page 6, paragraph [0089]).

As per claim 24, Huang discloses each entry of the virtual list comprising a node status and the management node further operable to change the status of each of the subset of nodes to "available" (col. 10, lines 50-56).

As per claim 25, Huang discloses the management node being further operable to: determine dimensions of the terminated job based on one or more job parameters and an associated policy; dynamically allocate a second subset of the plurality of nodes to the terminated job based on the determined dimensions (col. 17, lines 1-21) executing the terminated job on the allocated second subset (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 26, Huang discloses the second subset comprising a substantially similar set of nodes to the first subset (Fig. 2).

As per claim 27, Huang discloses the management node being operable to dynamically allocate the second subset comprises the management node being operable to: determine an optimum subset of nodes from a topology of unallocated nodes; and allocate the optimum subset (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 28, Huang discloses the management node being further operable to: locate a replacement node for the failed node; and update the logical entry of the

failed node with information on the replacement node (col. 7, lines 5-67 through col. 8, lines 1-9).

As per claim 29, Huang discloses the management node being operable to determine that one of the plurality of nodes has failed comprises the management node operable to determine that a repeating communication has not been received from the failed node (col. 17, lines 21-30).

As per claim 30, Huang discloses the management node is operable to determine through polling one of the plurality of nodes has failed (col. 8, lines 43-63).

Allowable Subject Matter

The following is an examiner's statement of reasons for allowance:

After a complete search of all the relevant prior art the examiner has determined the claims are in condition for allowance. The following limitations when viewed in combination with the remainder of the claim as a whole place this application in condition for allowance.

As per claim 31, the examiner finds the novel and non obvious feature of claim 31, when read as whole to be a first switch integrated to the first card, the first processors communicably coupled to the first switch, the first switch operable to communicably couple the first processors to at least six second cards each comprising at least two second processors integrated to the second card and a second switch

integrated to the second card operable to communicably couple the second processors to the first card and at least five third cards each comprising at least two third processors integrated to the third card and a third switch integrated, to the third card; the first processors being operable to communicate with particular second processors on a particular second card via the first switch and the second switch on the particular second, card; the first processors being operable to communicate with particular third processors on a particular third card via the first switch, a particular second switch on a particular second card between the first card and the particular third card, and the third switch on the particular third card without communicating via either second processor on the particular second card.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

Response to Arguments

Applicant's arguments filed July 14, 2008 have been fully considered but they are not persuasive.

As per claims 1, 11, and 21 and in response to applicant's argument that Karpoff fails to disclose at lease two processors integrated to the card and a switching fabric integrated to the same card, the Examiner respectfully disagrees and would like to point out that Karpoff discloses each node comprising a switching fabric integrated to a card

and at least two processors integrated to the card (FIG. 4A, a typical INFINIBAND® Architecture 20 includes one or more Central Processing Units (CPUs) 30, a Memory Controller 28, a Host Interconnect 29, a Host Channel Adapter (HCA) 22, a Target Channel Adapter (TCA) 24, and one or more Switches 26 ([0089]). Karpoff further discloses nodes attached to the fabric can be assembled into logical subsets or partitions in order to group hosts or devices with like attributes, much like zoning capabilities of Fiber Channel fabrics ([0087]).

Further Karpoff discloses implementing the RAID engine into the controller and placing cache on the blades within the switching fabric ([0041]). Karpoff discloses many implementations are possible with respect to the switch blade ([0104], [0105], and [0111]).

Conclusion

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Elmira Mehrmanesh whose telephone number is (571) 272-5531. The examiner can normally be reached on 9-5 M-F.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Robert W. Beausoliel can be reached on (571) 272-3645. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/Robert W. Beausoliel, Jr./

Supervisory Patent Examiner, Art Unit 2113